CLAIMS

1. A method for a coordinated bringup of a repaired storage appliance in a storage appliance cluster, the repaired storage appliance having a disk subsystem, the method comprising the steps of:

asserting a GIVEWAIT state in a predetermined memory location of the repaired storage appliance;

releasing disk reservations in response to detection of the asserted GIVEWAIT state by a surviving storage appliance;

initializing the disk subsystem of the repaired storage appliance;

asserting a MBWAIT state in the predetermined memory location; and

performing a giveback operation by the surviving storage appliance in response to detecting the MBWAIT state.

2. The method of claim 1 further comprising the steps of:

completing the repaired storage appliance initialization; and

processing data access requests by the repaired storage appliance.

3. The method of claim 1 wherein the predetermined memory location comprises a state data structure within a memory of the repaired storage appliance.

4. The method of claim 1 wherein the surviving storage appliance detects the GIVEWAIT state by performing a remote direct memory access read operation to the predetermined memory location.

5. The method of claim 1 wherein the surviving storage appliance detects the MBWAIT state by performing a remote direct memory access operation of the predetermined memory location.

6. The method of claim 1 wherein the surviving storage appliance ceases to process data access requests directed to the repaired storage appliance after performing the give-back operation.

7. A storage appliance for use in a storage system cluster, the storage appliance comprising:

a storage operating system having a cluster failover layer adapted to perform a coordinated bringup operation in association with a partner storage appliance, wherein the coordinated bringup operation comprises the steps of:

(i) asserting a first state in a predetermined memory location of the storage appliance;

(ii) initializing a disk subsystem of the repaired storage appliance in response to detecting a release of disk reservations by a partner storage appliance;

(iii) asserting a second state in the predetermined memory location;

(iv) processing data access requests directed to the storage appliance after a giveback operation performed by the partner storage appliance; and

whereby a period of time during which clients of the storage system are without connectivity is minimized.

8. The storage appliance of claim 6 wherein the cluster failover layer is further adapted to perform routine remote direct and memory access read operations to the partner storage appliance to detect a state of the partner storage appliance.

9. The storage appliance of claim 8 wherein the second state comprises a MBWAIT state.

10. The storage appliance of claim 8 wherein the first state comprises a GIVEWAIT state.

11. A method for a coordinated bringup of a repaired storage appliance in a storage appliance cluster, the repaired storage appliance having a disk subsystem, the method comprising the steps of:

asserting a first state in a predetermined memory location of the repaired storage appliance;

releasing disk reservations in response to detection of the asserted first state by a surviving storage appliance;

initializing the disk subsystem of the repaired storage appliance;

asserting a second state in the predetermined memory location; and

performing a giveback operation by the surviving storage appliance in response to detecting the second state.

12. The method of claim 11 wherein the predetermined memory location comprises a state data structure within a memory of the repaired storage appliance.

13. The method of claim 11 wherein the surviving storage appliance detects the first state by performing a remote direct memory access read operation to the predetermined memory location.

14. The method of claim 11 wherein the surviving storage appliance detects the second state by performing a remote direct memory access operation of the predetermined memory location.

15. The method of claim 11 wherein the surviving storage appliance ceases to process data access requests directed to the repaired storage appliance after performing the giveback operation.

16. The method of claim 11 wherein the first state comprises a GIVEWAIT state.

17. The method of claim 11 wherein the second state comprises a MBWAIT state.

18. The method of claim 11 wherein the set of disk reservations comprises small computer systems interface reservations.

19. A computer readable medium, including program instructions executing on a storage appliance, for a coordinated bringup of a repaired storage appliance in a storage appliance cluster, the repaired storage appliance having a disk subsystem, the computer readable medium including instructions for performing the steps of:

asserting a GIVEWAIT state in a predetermined memory location of the repaired storage appliance;

releasing disk reservations in response to detection of the asserted GIVEWAIT state by a surviving storage appliance;

initializing the disk subsystem of the repaired storage appliance;

asserting a MBWAIT state in the predetermined memory location; and

performing a giveback operation by the surviving storage appliance in response to detecting the MBWAIT state.

20. The computer readable medium of claim 19 further comprising the steps of:

completing the repaired storage appliance initialization; and

processing data access requests by the repaired storage appliance.

21. The computer readable medium of claim 19 wherein the predetermined memory location comprises a state data structure within a memory of the repaired storage appliance.

22. The computer readable medium of claim 19 wherein the surviving storage appliance detects the GIVEWAIT state by performing a remote direct memory access read operation to the predetermined memory location.

23. The computer readable medium of claim 19 wherein the surviving storage appliance detects the MBWAIT state by performing a remote direct memory access operation of the predetermined memory location.